arXiv:1409.0761v1 [q-bio.PE] 2 Sep 2014

AN ALGORITHM FOR CONSTRUCTING PRINCIPAL GEODESICS IN PHYLOGENETIC TREESPACE

TOM M. W. NYE

Abstract. Most phylogenetic analyses result in a sample of trees, but summarizing and visualizing these samples can be challenging. Consensus trees often provide limited information about a sample, and so methods such as consensus networks, clustering and multidimensional scaling have been developed and applied to tree samples. This paper describes a stochastic algorithm for constructing a principal geodesic or line through treespace which is analogous to the first principal component in standard Principal Components Analysis. A principal geodesic summarizes the most variable features of a sample of trees, in terms of both tree topology and branch lengths, and it can be visualized as an animation of smoothly changing trees. The algorithm performs a stochastic search through parameter space for a geodesic which minimises the sum of squared projected distances of the data points. This procedure aims to identify the globally optimal principal geodesic, though convergence to locally optimal geodesics is possible. The methodology is illustrated by constructing principal geodesics for experimental and simulated data sets, demonstrating the insight into samples of trees that can be gained and how the method improves on a previously published approach. A java package called GeoPhytter for constructing and visualising principal geodesics is freely available from `www.ncl.ac.uk/~ntmwn/geophytter`.

This is a postprint of an article published in IEEE Transactions in Computational Biology and Bioinformatics

©IEEE/ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in IEEE Transactions in Computational Biology and Bioinformatics, Vol. 11, March-April 2014, `http://doi.ieeecomputersociety.org/10.1109/TCBB.2014.2309599`

1. Introduction

Samples of phylogenetic trees arise in many different contexts in phylogenetics: examples include Bayesian posterior samples, bootstrap

Date: May 2014.

samples and collections of trees from different genetic loci. Understanding and summarizing the information present in these samples is challenging due to the difficulty of representing and visualizing regions of the space of possible trees. Consensus trees are typically used to summarize samples. These represent the features on which trees in the sample tend to agree, and each edge is usually labelled with the proportion of trees in the sample containing that edge. Many different consensus methods have been proposed, differing in the particular features considered and the way conflicts between trees are resolved [1]. For example, the Adams consensus tree [2, 3] is constructed using the relationship between every triplet of leaves represented in a sample of rooted trees, and it does not take branch lengths on the trees into account. Alternatively, the average consensus tree [4] is defined using the matrix of path length distances between pairs of leaves on each tree, and it therefore depends on branch lengths. By definition, consensus trees only reflect limited aspects of the input sample. Specifically, they do not indicate alternative tree topologies; they cannot capture correlations between different features; and many do not incorporate information about variability in branch lengths. Consensus networks [5] were developed to address the first of these issues. They represent conflicting phylogenetic signals from different trees in a single network structure. Standard tools from multivariate data analysis, such as clustering [6, 7] and multidimensional scaling [8] have also been adapted and applied to samples of trees as a means of representing information in the samples more accessibly.

This paper presents a complementary approach which is analogous to a form of principal components analysis (PCA) adapted to the geometry of the space of phylogenetic trees. Given a sample of different phylogenies sharing the same set of taxa, we construct a principal geodesic, or line, in treespace which is a 'best fit' to the data in a well-defined sense. The principal geodesic can be thought of as a 1-parameter continuous family of trees which summarizes the sample, as opposed to the point summary provided by a consensus tree. In common with consensus network methods, the principal geodesic can represent a collection of alternative tree topologies, depending on the input sample. By analogy with the first principal component in standard PCA, the principal geodesic represents the most variable features of the sample, in terms of both topology and branch length. As for standard PCA, a 'proportion of variance' summary statistic can be computed for each principal geodesic. This measures the amount of variability captured by the principal geodesic in relation to the degree of scatter around the geodesic, and therefore indicates how well the principal geodesic summarizes the sample.

Our approach relies heavily on geometrical properties of the space of phylogenetic trees. The space of all possible phylogenetic trees on a fixed set of taxa forms a geometric space which we refer to as *treespace* [9]. The space is equipped with a metric, usually called the *geodesic metric*, and any pair of points can be joined by a unique *geodesic* i.e. a path of minimal length. A variety of other metrics exist for measuring differences between phylogenetic trees. However, the geodesic metric and the picture of treespace provided by Billera et al [9] provide the geometry necessary to perform an analog of PCA: namely, a notion of a straight line and a well-defined projection operation onto any given line. Other metrics do not provide these geometrical elements.

A sample of trees can be regarded as being drawn from some distribution on treespace, and our approach attempts to characterize this distribution by approximating the sample with a principal geodesic. An important feature of the approach is that it uses branch length information as well as topological information from the input sample of trees, since information about the shape of the distribution in treespace would be lost by marginalising out the branch lengths. For example, if most of the sampled trees had the same topology, then variability in the sample would largely be comprised of variability in the branch lengths for that topology. Our approach can characterise the variability equally well in this situation, or conversely for samples which are widely dispersed over different topologies, and it therefore offers a unified approach independent of the nature of the sample. Branch lengths and tree topology are intimately related since (i) tree topology can be changed by continuously shrinking edges to length zero and expanding out alternative edges, and (ii) the relative branch lengths in a species tree affect the inferred topology in various ways when constructing a phylogeny from genetic data. For example, short internal edges can be harder to infer and so lead to alternative topologies. Although biologists are often primarily concerned with the different tree topologies represented within a sample, incorporating branch lengths leads to a more complete picture.

Although normally applied to data in vector spaces, PCA can be adapted to work in other geometrical spaces. Most importantly, PCA has been reformulated in terms of geodesic geometry on Riemannian manifolds [10, 11] including Lie groups and shape spaces [12, 13]. These ideas have also been extended to certain spaces of trees [14]. Our algorithm is based on the same type of reformulation, and the first step is to re-express standard PCA in terms of the following schema. Suppose $x_1, x_2, \ldots, x_n$ is a set of points in a vector space equipped with an inner product and induced metric $d(\cdot, \cdot)$. The zero-th order principal component can be defined as the point θ_0 which minimizes the sum of squared distances $\sum d(x_i, \theta_0)^2$. Similarly, the first order principal component θ_1

is the line which minimizes $\sum d(x_i, P_1 x_i)^2$ where P_1 denotes projection onto θ_1. This definition extends to the k-th order principal component θ_k which is the k-dimensional subspace which minimizes the sum of squared projected distances $\sum d(x_i, P_k x_i)^2$ where P_k is projection onto θ_k. Algebraic solutions can be obtained for these minimization problems for PCA in vector spaces. In particular θ_0 is just the sample mean $\bar{x}$, and θ_1 is the line through $\bar{x}$ spanned by the eigenvector of the sample covariance matrix with largest eigenvalue. This schema for constructing θ_0 and θ_1 applies exactly as stated above to more general metric spaces equipped with some notion of lines and projection onto lines, such as treespace. In terms of treespace geometry, $x_1, \ldots, x_n$ is a sample of trees, θ_1 is a geodesic in treespace and $d(\cdot, \cdot)$ is the geodesic metric. However, in the more general setting, algebraic solutions are not available and θ_0 is not generally a subspace of θ_1 [15]. This can be demonstrated by simple examples in treespace. In the context of general metric spaces, the zero-th order principal component θ_0 is more commonly known as the *Fréchet mean*, and a stochastic algorithm for computing the Fréchet mean in treespace has recently been developed [16, 17] based on previous work in more general metric spaces [18]. In this paper we present a stochastic algorithm for constructing θ_1 in treespace, or more precisely, a geodesic which minimises the sum of squared distances between points in the sample and their projections onto the geodesic. The algorithm searches through the set of geodesics for a optimal fit to the data, though convergence to local minima is possible. We call the algorithm *GeoPhytter*.

The present author has previously published an algorithm called ΦPCA for constructing the first principal component in treespace [19], and the methods presented here form an updated approach to the same problem, overcoming a number of limitations of the original algorithm. Most importantly, ΦPCA only considers geodesics which lie in a restricted class with simple geometrical properties as candidates for θ_1. The algorithm fails to construct the optimal geodesic for certain data sets on account of this restriction, and in some cases ΦPCA fails to produce any output at all, such when all the input trees have the same topology. Later in the paper we give examples of biological data sets on which ΦPCA does not identify an optimal principal geodesic, and give a more thorough technical explanation of the improvements GeoPhytter represents. Other authors [20] have more recently also described an algorithm for constructing an approximate principal geodesic in treespace, for which the geodesic is constrained to lie between points in the data set. We apply this algorithm to experimental data sets later in the paper to compare with results obtained using GeoPhytter.

The remainder of the paper has the following structure. After reviewing the geometry of phylogenetic treespace, we then present the

stochastic algorithm for constructing principal geodesics. Detailed results are given for two data sets: (i) a sample of gene trees for a set of archaea for which gene conversion has affected certain loci and (ii) a simulated bootstrap sample of trees affected by long branch attraction. These illustrate potential biological applications of principal geodesic analysis in treespace. We also give very brief results for some other data sets to indicate the type of characterization of tree samples that the analysis can provide. Finally, we conclude with a discussion of computational issues and possibilities for further research.

2. Methods

2.1. **Geometry in treespace.** We need to describe the geometry of phylogenetic treespace in order to specify our algorithm fully. More details about the structure and geometry of treespace are given by Billera et al [9], and the account presented here is brief. A phylogenetic tree represents the evolutionary relationships between a set of objects, called *taxa*, which label the leaves of the tree. Treespace $\mathcal{T}_N$ is the set of all unrooted trees on the set of taxa $S = \{1, 2, \ldots, N\}$. The trees have positive-valued edge weights or *branch lengths*. Sometimes it is convenient to ignore the branch lengths on a tree, in which case we obtain the tree *topology*. Topologies are referred to as *resolved* if all the vertices apart from the leaves have degree 3, but are otherwise called *unresolved*. A *split* is a bipartition of S, and every edge in an unrooted tree is associated with the split of taxa induced by cutting the edge. Splits are often written in the notation $abc\ldots|xyz\ldots$ where $\{a, b, c, \ldots\} \cup \{x, y, z, \ldots\}$ is a disjoint union of S. Although we work throughout with unrooted trees, the space of rooted trees can be obtained by adding in an additional taxon representing the ancestor of all the taxa in S, and so the space of rooted trees is essentially isomorphic to $\mathcal{T}_{N+1}$. Each tree in $\mathcal{T}_N$ contains up to $N-3$ internal edges and exactly N *pendant* edges, i.e. those which contain a leaf. Treespace consists of different regions, each corresponding to the set of trees with a particular fully-resolved topology. If we consider a single fully-resolved tree topology τ, and ignore the lengths of the pendant edges, then by associating each internal edge with a coordinate axis in $\mathbb{R}^{N-3}$ there is a bijection between the set of trees with topology τ and the interior of the *positive orthant* $\mathcal{O}_\tau = \mathbb{R}_+^{N-3}$ where $\mathbb{R}_+ = \{x \in \mathbb{R} : x \geq 0\}$. The faces of $\mathcal{O}_\tau$ correspond to unresolved trees. Continuing to ignore pendant edge lengths, it follows that as a set, $\mathcal{T}_N$ corresponds to $\bigcup \mathcal{O}_\tau$, where the union is over all possible fully-resolved tree topologies.

The orthants corresponding to different topologies overlap along their faces in the following way. The codimension-1 faces of each orthant correspond to trees in which a single internal edge has contracted to length zero, resulting in a vertex with degree 4. Such a vertex can be resolved

in three different ways leading to two new topologies τ', τ'' as well as the original τ. It follows that each codimension-1 face of $\mathcal{O}_\tau$ is identified with corresponding faces in $\mathcal{O}_{\tau'}$ and $\mathcal{O}_{\tau''}$. The topological operation corresponding to shrinking down an edge and replacing it with one of its two alternatives is referred to as *nearest neighbor interchange* (NNI) and adjacent regions in $\mathcal{T}_N$ contain trees related by NNI. Higher codimension faces similarly form the intersection of multiple orthants, and the origin of tree space, corresponding to a tree with no internal edges (or *star-tree*), is the unique point in the intersection of all the orthants. Adding back in the pendant edges, $\mathcal{T}_N$ is the product $\mathbb{R}_+^N \times \bigcup \mathcal{O}_\tau$ where each point in $\mathbb{R}_+^N$ determines a set of pendant edge-lengths.

Insight into the structure of treespace can be gained by considering low-dimensional examples. For $N = 4$ taxa, there are 3 possible unrooted tree topologies, so (ignoring pendant edges for simplicity) $\mathcal{T}_4$ consists of three copies of $\mathbb{R}_+$ joined at together at a point which corresponds to a star-tree. This is illustrated in Fig. 1. For $N = 5$ there are 15 unrooted tree topologies, so $\mathcal{T}_5$ is formed from 15 orthants each of which is a copy of $\mathbb{R}_+^2$. As Fig. 2 shows, each orthant is joined to two others along each codimension-1 face, corresponding to the two choices for nearest-neighbor interchange on a particular edge of a tree. The combinatorial structure of the orthants corresponds to a copy of the Petersen graph, as illustrated by in Fig. 2, and $\mathcal{T}_5$ is the topological cone of this graph. For $N = 6$ the structure is more complicated again, since unlike the case for $N = 4$ and 5, there are different possible unlabelled tree shapes. However, the underlying principles remain the same for all values of N: each orthant corresponds to all the trees with a particular topology, and every orthant is joined to two others along each codimension-1 face.

The metric structure on $\mathcal{T}_N$ is obtained as follows. Any two trees x_1, x_2 with the same topology τ can be joined by a straight line segment in $\mathbb{R}_+^N \times \mathcal{O}_\tau$ and $d(x_1, x_2)$ is defined as the standard Euclidean (L^2) length of that segment. If $x_1, x_2 \in \mathcal{T}_N$ have different topologies then $d(x_1, x_2)$ is defined as the length of the shortest path between x_1 and x_2 which consists of straight line segments in each orthant, where path length is defined as the sum of individual segment lengths. Billera et al[9] showed that this determines a well-defined metric $d(\cdot, \cdot)$ on $\mathcal{T}_N$ and that the shortest-length path, or *geodesic*, between any two points is unique. A $O(N^4)$ algorithm has been developed for constructing geodesics [21]: its input is a pair of trees x_1, x_2 and the algorithm outputs a sequence of straight line segments contained in a sequence of orthants linking x_1 to x_2. Any pair of trees x_1, x_2 with different topologies τ_1, τ_2 can be connected by a path consisting of the line segment in $\mathcal{O}_{\tau_1}$ joining x_1 to the origin and the segment in $\mathcal{O}_{\tau_2}$ joining the origin

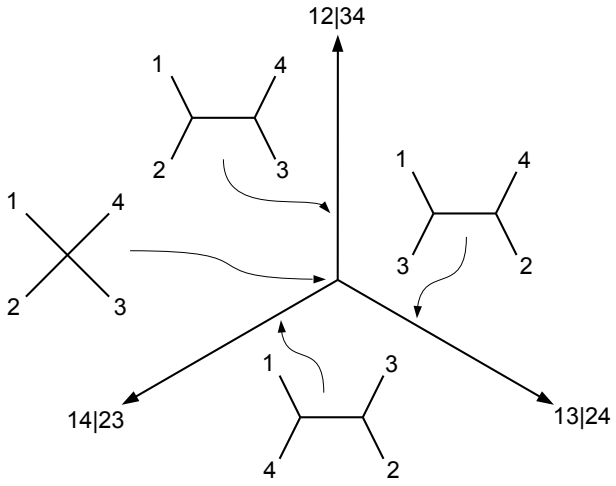

FIGURE 1. Treespace for $N = 4$ taxa consists of three copies of $\mathbb{R}_+$, joined together at a point which corresponds to a star tree. Each copy of $\mathbb{R}_+$ is labelled with its corresponding split, and the position point along $\mathbb{R}_+$ determines the length of the edge associated with the split.

to x_2. Such paths are called *cone paths* and they sometimes coincide with the geodesic between x_1 and x_2.

We use the algorithm of Owen and Provan[21] to construct geodesics between points in treespace. Given a point $x \in \mathcal{T}_N$ and a geodesic γ between two points $\gamma_1, \gamma_2 \in \mathcal{T}_N$ there is a unique closest point on γ to x, which we call the *projection* of x, denoted $P_\gamma(x)$. Sometimes x is closest to an end-point of γ in which case $P_\gamma(x) = \gamma_i$ for $i = 1$ or 2. The existence of the projection is a result of the particular mathematical structure of $\mathcal{T}_N$ [9], and an efficient algorithm for calculating the projection $P_\gamma(x)$ was presented in [19]. We use the same algorithm to perform projections here. The algorithm involves a path-length parametrization $\gamma(s)$ of γ, and uses a golden ratio search algorithm to find the value of s which minimises $d(\gamma(s), x)$.

In contrast to standard PCA, in which the first principal component is an infinite line, in treespace it is advantageous to restrict attention to geodesic segments of finite length. To illustrate this, consider a situation in which the projected points $P_\gamma(x_i)$ all lie in a single orthant $\mathcal{O}_\tau$ when γ has been chosen to minimise the sum of perpendicular squared distances. In this case, the extension of γ into other orthants is entirely arbitrary, and the data are only informative about the restriction of geodesics to $\mathcal{O}_\tau$. This situation would arise when the data

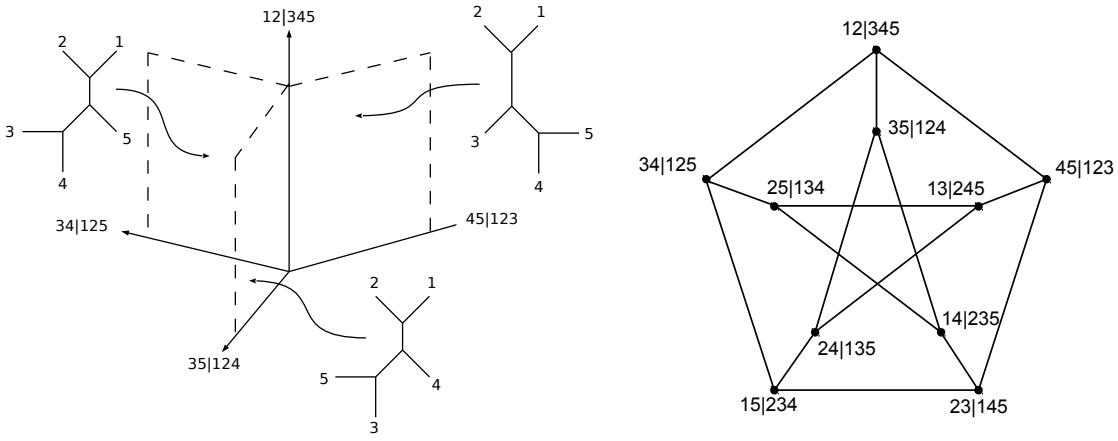

FIGURE 2. Treespace for $N = 5$ taxa. Left: three orthants in $\mathcal{T}_5$ and their corresponding tree topologies. The position within an orthant determines the two internal edge lengths on each tree, and every orthant is glued to two others along each codimension-1 face. Right: the combinatorial structure of the orthants corresponds to the Petersen graph. Each vertex on the graph corresponds to a codimension-1 face of an orthant, and is labelled with the split assigned length zero on that face. The edges on the graph correspond to the 15 orthants in $\mathcal{T}_5$. The three edges connected to the vertex labelled 12|345 correspond to the orthants shown on the left. $\mathcal{T}_5$ is the topological cone of this graph.

consist of a tightly clustered collection of trees all having the same topology. It follows that more generally, arbitrary samples of trees will not always determine a unique infinite principal geodesic, and so we consider geodesic segments of finite length. The suitability of finite geodesic segments as opposed to infinite geodesics for principal geodesic analysis has previously been recognized in [20].

2.2. **Algorithm.** Given a sample of trees $x_1, x_2, \ldots, x_n$, we wish to find a geodesic $\gamma = (\gamma_1, \gamma_2)$ which minimises the objective function:

$$\omega(\gamma; x_1, \ldots, x_n) = \sum_i d(x_i, P_\gamma(x_i))^2 + \min_i d(\gamma_1, P_\gamma(x_i))^2 + \min_i d(\gamma_2, P_\gamma(x_i))^2. \quad (1)$$

The first term is the sum of squared perpendicular distances, while the last two terms are included so that the global minimum is a geodesic of minimum length. The last two terms only contribute to the objective when all the data points project onto the interior of γ. The projection

of $x_1, \ldots, x_n$ onto γ can be parallelized trivially by partitioning the data set, and our implementation takes advantage of parallel architectures. At each step of the algorithm we maintain two points $\gamma_1, \gamma_2 \in \mathcal{T}_N$ and the geodesic $\gamma = (\gamma_1, \gamma_2)$. We use the notation $\bar{\gamma}_1 = \gamma_2$ and $\bar{\gamma}_2 = \gamma_1$ so that $(\bar{\gamma}_1, \bar{\gamma}_2)$ represents the same geodesic but with the reverse orientation. The algorithm is initialized with γ_1, γ_2 taken to be a random pair of points in the data set. The algorithm uses a set of stochastic 'moves' f_k, $k \in K$ which randomly perturb points in treespace. Five different moves are available, so $K \subset \{1, 2, 3, 4, 5\}$, and the moves are specified below. The algorithm operates by repeating the following procedure until convergence is obtained.

for $k \in K$ **do**
  for $i = 1, 2$ **do**
    1) Generate a random tree $y = f_k(\gamma_i)$ and construct the geodesic $\eta = (\gamma_i, y)$.
    2) Consider the set of geodesics $\{(\bar{\gamma}_i, z) : z \in \eta\}$. Use a 1-dimensional optimization method (golden ratio search) to find $z \in \eta$ which minimises the objective ω on the restriction to this set.
    3) Set $\gamma_i = z$.
    4) Either extend or contract γ to deal with projection onto the end-points, as detailed below.
    5) The resulting geodesic γ is used as the starting point for the next iteration, at step 1 above.
  end for
end for

An illustration of the procedure is given in Fig. 3. The moves $f_1, f_2, \ldots, f_5$ for randomly perturbing the end-points of γ are specified as follows.

Gaussian random walk: Given an initial fully-resolved tree $y_0 = x$ we will define a random sequence of trees $y_1, y_2, \ldots, y_{n_{\mathrm{RW}}}$ for some fixed number of steps n_{RW} and take $f_1(x) = y_{n_{\mathrm{RW}}}$. To define this sequence, we fixing an ordering of the internal edges $e_1, \ldots, e_{N-3}$ in x and let $l(e) \in \mathbb{R}^+$ denote the length of e. To obtain y_{j+1} from y_j we consider each internal edge e in y_j in turn and apply the follow procedure:

(1) Generate a normally distributed random variate $z \sim N(0, \sigma^2)$ and let $l^* = l(e) + z$.
(2) If $l^* \geq 0$ set the length of e to be l^*.
(3) Conversely if $l^* < 0$ choose one of the two NNI replacements for e uniformly at random, denote it e', and replace e in y_j by e', setting $l(e') = |l^*|$.

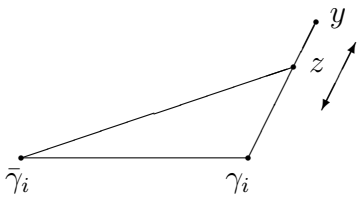

FIGURE 3. Illustration of one step of the algorithm. A point y is obtained by applying a stochastic rule f_k to one end γ_i of γ, and the geodesic $\eta = (\gamma_i, y)$ is constructed. The objective ω is evaluated for different geodesics $(\bar{\gamma}_i, z)$ where $z \in \eta$ and $\bar{\gamma}_i$ is the opposite end to γ_i. Each evaluation involves projecting the data points onto the candidate geodesic $(\bar{\gamma}_i, z)$. By applying a golden ratio search method, γ_i is replaced by the value of z which minimises the objective.

We call this a *Gaussian random walk* on treespace, and it provides a straightforward means of randomly perturbing trees. For the results in this paper, a fixed number of steps $n_{\mathrm{RW}} = 10$ was used for the Gaussian random walk, and σ^2 was fixed to give a certain approximate probability that each edge was replaced by at least one NNI during the walk. If the mean internal edge length in the data set is denoted $\bar{l}$, then the probability of at least one NNI on a given edge is approximately $\Pr(X < 0)$ where $X \sim N(\bar{l}, n_{\mathrm{RW}}\sigma^2)$. We took $\sigma^2 = \bar{l}/0.85n_{\mathrm{RW}}$, to give an approximate probability of 20%.

NNI: $f_2(x)$ is obtained by sampling an internal edge e in x uniformly at random, and randomly selecting an NNI replacement e' for e from the two possibilities. The length of e' in $f_2(x)$ is set to a fixed value (the largest edge length observed in the data set).

Random data point: $f_3(x)$ is taken to be a data point x_i sampled uniformly at random from $x_1, \ldots, x_n$.

End-point move: Suppose that data points $x_{r_1}, x_{r_2}, \ldots, x_{r_D}$ project onto the end γ_i of the geodesic. An approximate Fréchet mean is computed for these data points via the algorithm of Miller et al [17] and Bacak [16], as follows. First $x_{r_1}, x_{r_2}, \ldots, x_{r_D}$ are randomly permuted, and then we fix $y_1 = x_{r_1}$. To obtain y_{j+1} from y_j we construct the geodesic $(y_j, x_{r_{j+1}})$ and let y_{j+1} be the point a proportion $1/(j+1)$ along the geodesic. This rule is applied iteratively for $j = 2, 3, \ldots, D$ to obtain $f_4(\gamma_i) = y_D$. If no points project onto γ_i then this move is not performed.

Pendant edge lengths: $f_5(x)$ is obtained by randomly perturbing the pendant edge lengths in x. Specifically, for each pendant edge

e with length $l(e)$, e is assigned a new length drawn from a gamma distribution with mean $l(e)$ and variance ν^2. We took $\nu^2 = n_{\mathrm{RW}}\sigma^2$ so that this move had comparable variance to the Gaussian random walk.

The software allows the user to apply the algorithm with any subset $K \subset \{1, \ldots, 5\}$ of moves rather than all five, and also specify the exact order in which the moves are applied. Move f_5 is only required if pendant edges are included in the analysis. By default the software uses moves 1–3 applied in index order, and pendant edges are ignored – the metric is computed without including them. However, a software option allows the user to include pendant edges in the analysis together with move f_5.

2.3. **Contracting and extending γ.** Every time γ is updated at step 4 in the algorithm above, γ is either extended or contracted in order to reduce the objective $\omega(\gamma)$ further, in the following way. Consider the situation when none of the data points project onto a particular end of γ. Without loss of generality, suppose this end is γ_1. Then γ_1 can be replaced with the closest projected point $P_\gamma x_i$, thereby making the middle term in (1) vanish but leaving the other terms unchanged. Thus contraction always reduces the objective whenever no points project onto either end of γ. The contraction operation leads to estimated principal geodesics which are of minimal length, and it ensures that at least one data point projects onto each end of γ at each step of the algorithm.

Conversely, if any data points project onto an end γ_i and that end lies in the interior of an orthant, then γ can be extended up to the boundary of the orthant and the projection of these data points re-calculated. For example, if γ_1 lies in the interior of $\mathbb{R}^N_+ \times \mathcal{O}_{\tau_1}$ then since $\gamma \cap (\mathbb{R}^N_+ \times \mathcal{O}_{\tau_1})$ is a line segment it has a unique extension either from γ_1 up to the boundary of $\mathbb{R}^N_+ \times \mathcal{O}_{\tau_1}$ or an extension out to infinity. Extension across boundaries is non-unique, since there is always a choice about which adjacent orthant to extend into, and so we only make use of the unique extension to the boundary. Now suppose without loss of generality that a data point x projects onto the end γ_1 and define $\phi(s) = d(x, \gamma(s))$ where $\gamma(s)$ is a linear parametrization of γ for $s \in [0, 1]$. The extension of γ in $\mathcal{O}_{\tau_1}$ enables the domain of ϕ to be extended to an open neighbourhood $s \in (-\epsilon, \epsilon)$ of zero, and the definition of the metric ensures ϕ is continuous on this domain. Since x projects onto the point $s = 0$, ϕ decreases as s tends to zero from above. It follows that extension and re-calculation of the projection of x will lead to a reduction in the objective except in the special case that ϕ has a minimum at $s = 0$. Note that the shortest possible extension which improves the objective is taken, in the sense that after extension the last two terms of (1) vanish. If many points project onto an end-point, and

that end-point lies on the face of an orthant (i.e. the tree is unresolved) it might indicate that a geodesic with a better objective value could be obtained by moving the end-point into a neighbouring orthant, and that consequently the geodesic is sub-optimal. However, this seems to occur rarely in practice.

2.4. **Convergence and local minima.** The algorithm is not guaranteed to converge to a global minimum. In practice the optimization procedure is halted if the fractional change in objective is below a fixed threshold for a large number of iterations dependent on the number of trees in the data set and the number of taxa. To test optimality of the resulting geodesic, the algorithm can be run several times, with a different starting value for γ each time. If multiple different runs converge to very similar geodesics, that suggests a global optimium has been found, though it cannot be guaranteed. A unique global minimum for the objective may not exist: that is the case for standard PCA when several eigenvectors of the data covariance matrix share the same maximimal eigenvalue. Since the analysis in treespace essentially reduces to standard PCA when all the data points are highly concentrated in a single orthant, treespace also lacks unique optima in general.

2.5. **Rationale.** The high dimensionality and complex combinatorial structure of treespace motivated the decision to develop a stochastic algorithm to optimize the objective function in equation (1). A simple Localised Random Search algorithm [22] would have the following form:

(1) Start with a random geodesic γ.
(2) Randomly perturb γ, for example by performing a Gaussian random walk on one of its ends, to obtain a new geodesic γ^*.
(3) If $\omega(\gamma^*) < \omega(\gamma)$ set $\gamma = \gamma^*$.
(4) Repeat from step 2, unless some convergence criterion has been satisfied.

This algorithm is simple to implement, but its performance depends on the exact nature of stochastic innovation at step 2. In general, if an innovation with small variance is used, then the proposed improvement γ^* tends to be close to γ, and the algorithm slowly moves downhill to settle at a local minimum. Conversely, by using an innovation with large variance, larger steps are possible enabling the algorithm to traverse between local minima more easily, but at step 3 the proposed improvement γ^* is often rejected. We tested a version of Localised Random Search for which the innovation consisted of a Gaussian random walk on a randomly chosen end of the geodesic γ. This algorithm performed very badly for a range of different variances in the random walk step, in particular since most proposals γ^* failed to improve the objective.

GeoPhytter is based on Localised Random Search, but it combines stochastic innovations with deterministic optimization over slices in treespace. The combination of the stochastic moves followed by deterministic optimization over slices gave much better performance in terms of computational time and convergence properties. In particular, it employs both 'local' innovations (like the NNI move) and 'global' innovations (namely the random data point move) which enable the algorithm to explore the local neighbourhood of γ while still having the potential to escape from local minima.

The specific moves $f_1, \ldots, f_5$ were designed under the following rationale. Moves f_1 and f_2 represent 'local' innovations (although the variance of the random walk can be adjusted to make it perform bigger steps), while f_3 enables larger 'global' moves. Additionally, f_1, f_2 enable γ to move out of the convex hull of the data set; f_3 cannot achieve this on its own, but it is easy to construct examples for which the principal geodesic does not lie in the convex hull of the data. Move f_4 tends to shift the ends of the segment γ outwards in the direction of the set of points projecting onto each end. The random perturbation of pendant edges, f_5, explores nearby geodesics for which the pendant edge lengths differ from γ. It provides a counterpart to f_1 and f_2, which do not affect the pendant edges.

The order in which the different moves were applied had little effect on the convergence properties of the algorithm. However, the convergence properties did depend on the selection of moves employed. Runs of the algorithm which used all the moves tended to converge faster than those using a limited set. Move f_3, selection of a random data point, is crucial: the algorithm performs relatively well using this move alone, but without it the algorithm converges more slowly and multiple runs often fail to converge to the same geodesic.

While pendant edges can optionally be included in any analysis, their inclusion tends to worsen convergence. In particular, when pendant edges are included the algorithm converges to local minima more frequently (indicated by multiple runs converging to different geodesics). A greater number of replicate runs is therefore required. Furthermore, for some data sets we analysed, the variability in pendant edge lengths swamped the signal from internal edges. Since biological interest principally lies with internal edges, we recommend that the software is used primarily for analyses ignoring pendant edges.

2.6. **Proportion of variance statistic.** In standard PCA the sum of squared distances from the sample mean decomposes as

$$\sum d(x_i, \bar{x})^2 = \sum d(x_i, P_1 x_i)^2 + \sum d(P_1 x_i, \bar{x})^2 \qquad (2)$$

where P_1 denotes projection onto the first principal component θ_1. The quantity $r^2 = \sum d(P_1 x_i, \bar{x})^2 / \sum d(x_i, \bar{x})^2$ is usually reported and is interpreted as the proportion of variance in the sample explained by θ_1. The Pythagorean theorem does not hold in treespace so a decomposition like (2) does not apply. Nonetheless, an analog of r^2 can be defined as follows. Let

$$d_\perp^2 = \sum d(x_i, P_\gamma(x_i))^2, \quad \text{and} \quad d_\parallel^2 = \sum |s_i - \bar{s}|^2$$

where $s_i = d(\gamma_1, P_\gamma(x_i))$ for $i = 1, \ldots, n$ and $\bar{s}$ is the mean of the s_i. Specifically s_i is the position along the geodesic γ of the projection of x_i obtained by the projection algorithm described in section 2.1. Equivalently, $d_\parallel^2$ can be defined as $\sum d(y_i, \bar{y})^2$ where $y_i = P_\gamma(x_i)$ and $\bar{y} = \gamma(\bar{s})$ is the Fréchet mean of $y_1, \ldots, y_n$. The quantity $r_\gamma^2 = d_\parallel^2/(d_\parallel^2 + d_\perp^2)$ is then the analog of r^2 in the standard analysis, and it is readily computed for any geodesic γ. For standard PCA in $\mathbb{R}^k$, r^2 is bounded below by $1/k$, with equality only in the case of isotropic data. This inequality does not hold in treespace, but it forms a useful baseline with which to interpret the statistics r_γ^2, by taking $k = N - 3$ (the number of internal edges). Note that a geodesic which minimizes $d_\perp^2$ does not necessarily maximize r_γ^2. This is essentially a result of the failure of the Pythagorean theorem in treespace, and examples based on trees containing 4 taxa can readily be constructed to demonstrate this.

2.7. **Comparison with existing algorithms.** The previously published algorithm, ΦPCA, was based on a competely different method for constructing principal geodesics, by building the principal geodesic up one split at a time. More importantly, ΦPCA made certain assumptions about the principal geodesic which we do not make here, as follows:

(1) In ΦPCA, the principal geodesic was forced to lie in a restricted set of geodesics, called *simple* geodesics. These have the property that as the geodesic is traversed, at most one edge is shrunk down to have zero length at a time. This rules out many possibilities, such as all geodesics which are cone paths, and the restriction was made for computational convenience rather than biological reasons. The algorithm presented here makes no such restrictions.
(2) ΦPCA sought to construct an infinite principal geodesic; we explained above why the data are only generally informative about a finite principal geodesic segment. On account of this, ΦPCA fails to find the principal geodesic in certain cases, for example when the data do not represent every topology that arises along some infinite extension of a finite principal geodesic.

(3) The principal geodesic in ΦPCA was forced to pass through a consensus tree under the assumption that this tree lay close to the Fréchet mean. The present algorithm does not restrict γ to pass through any mean or consensus tree. As explained in the introduction, the geodesic which minimises the sum of squared perpendicular distances does not necessarily contain the Fréchet mean.

The first of these points is probably the most important. Given these differences between the algorithms, it is easy to construct simulated data sets for which ΦPCA fails but GeoPhytter successfully identifies the principal geodesic (for example, constructing a data set by simulating trees distributed along a cone-path geodesic and then perturbing each tree slightly via a Gaussian random walk). In addition, there are biological reasons for believing that for some data sets the principal geodesic will not lie in the class of simple geodesics. For example, suppose that trees in some sample of gene trees have one of two topologies which are related by a sub-tree prune and regraft operation (SPR). This could arise if a subset of genes were horizontally transferred at some stage on the evolutionary tree. Depending on the particular topologies, geodesics between trees with topologies related by SPR are often cone paths, and so it is reasonable to anticipate that for some data sets, the principal geodesic would be a cone path and therefore lie outside the class of simple geodesics. The chaperonin data set analysed in the next section supports this intuition: the associated principal geodesics are not contained in the simple geodesic class.

A stochastic algorithm for constructing an approximate principal geodesic was also presented in [20]. The geodesic is constrained to lie between a pair of data points x_i, x_j. Pairs of data points are randomly sampled for a large number of iterations, and the pair which minimizes the sum of squared projected distances is taken as an approximate principal geodesic. We compare the performance of this algorithm with GeoPhytter in the next section.

3. Results

3.1. **Chaperonin data set.** We applied our method to a sample of trees constructed using an alignment of chaperonin gene sequences from archaea taken from a previous study. The original study [23] concerned duplication and gene conversion of chaperonin genes, and we analysed the same data set, kindly provided by the authors. An ancient gene duplication produced two copies of the chaperonin gene, α and β, in the archaea included in the study. The alignment used to construct phylogenies contained the α and β copies of the gene from 6 different archaea, giving a total of 12 sequences containing 1557 DNA sites. The

6 archaea were *Pyrodictium occultum*, *Aeropyrum pernix* and *Pyrobaculum aerophilum*, together with 3 closely related *Sulfolobus* species. The original analysis suggested two distinct topologies associated with different regions of the alignment. The first, which we will refer to as the *duplication topology*, separated the α from the β sequences into two highly supported clades, with very similar topologies within the two clades. The maximum likelihood tree inferred from the full alignment had this topology. In contrast, the second topology, which we will refer to as the *conversion topology*, had separate clades for three organisms (*P. occultum*, *A. pernix* and *P. aerophilum*), each clade containing the α and β sequences for the organism. This topology was associated with a subset of approximately 300 sites from three separate contiguous loci within a certain domain of the chaperonin protein. The topologies are shown in Fig. 4. Archibald and Roger [23] suggested that multiple independent gene conversion events had given rise to different regions in the gene supporting the two different topologies. We call the split separating the α sequences from the β sequences the *duplication* split, and conversely, call the split separating the *Sulfolobus* sequences from the others the *Sulfolobus* split.

We constructed phylogenetic trees in a similar way to the original paper, using exactly the same substitution model (general time-reversible plus gamma rate heterogeneity plus invariant sites). Maximum likelihood trees were inferred for a series of windows of the alignment each 200 nucleotides long, obtained by progressively sliding the window 10 nucleotides at a time. This gave a sample of 136 phylogenetic trees. The majority consensus of the sample resolved the *Sulfolobus* α and β sequences into two separate clades (both of which were present in the duplication and conversion topologies defined above), and contained the *Sulfolobus* split, but was otherwise unresolved. A principal geodesic was constructed using the algorithm presented in Sec. 2.2. Pendant edges were ignored in the analysis. Ten random starting points for the algorithm converged to the same principal geodesic segment, suggesting that the global optimum geodesic might have been found. Exactly one tree in the sample projected onto each end of the segment. Fig. 4 shows the two trees forming the extremities of the principal geodesic segment. The sums of squared distances associated with principal geodesic were $d_\perp^2 = 27.5$ and $d_\parallel^2 = 20.5$ so that $r_\gamma^2 = 43\%$, indicating that the principal geodesic represents a relatively large proportion of the variability within the sample.

The principal geodesic is associated with a transition between a topology containing the duplication split and a topology containing the *Sulfolobus* split. It essentially represents variability in the relationship between the *Sulfolobus* α and β clades. The two ends of the principal geodesic do not exactly correspond to the duplication and conversion

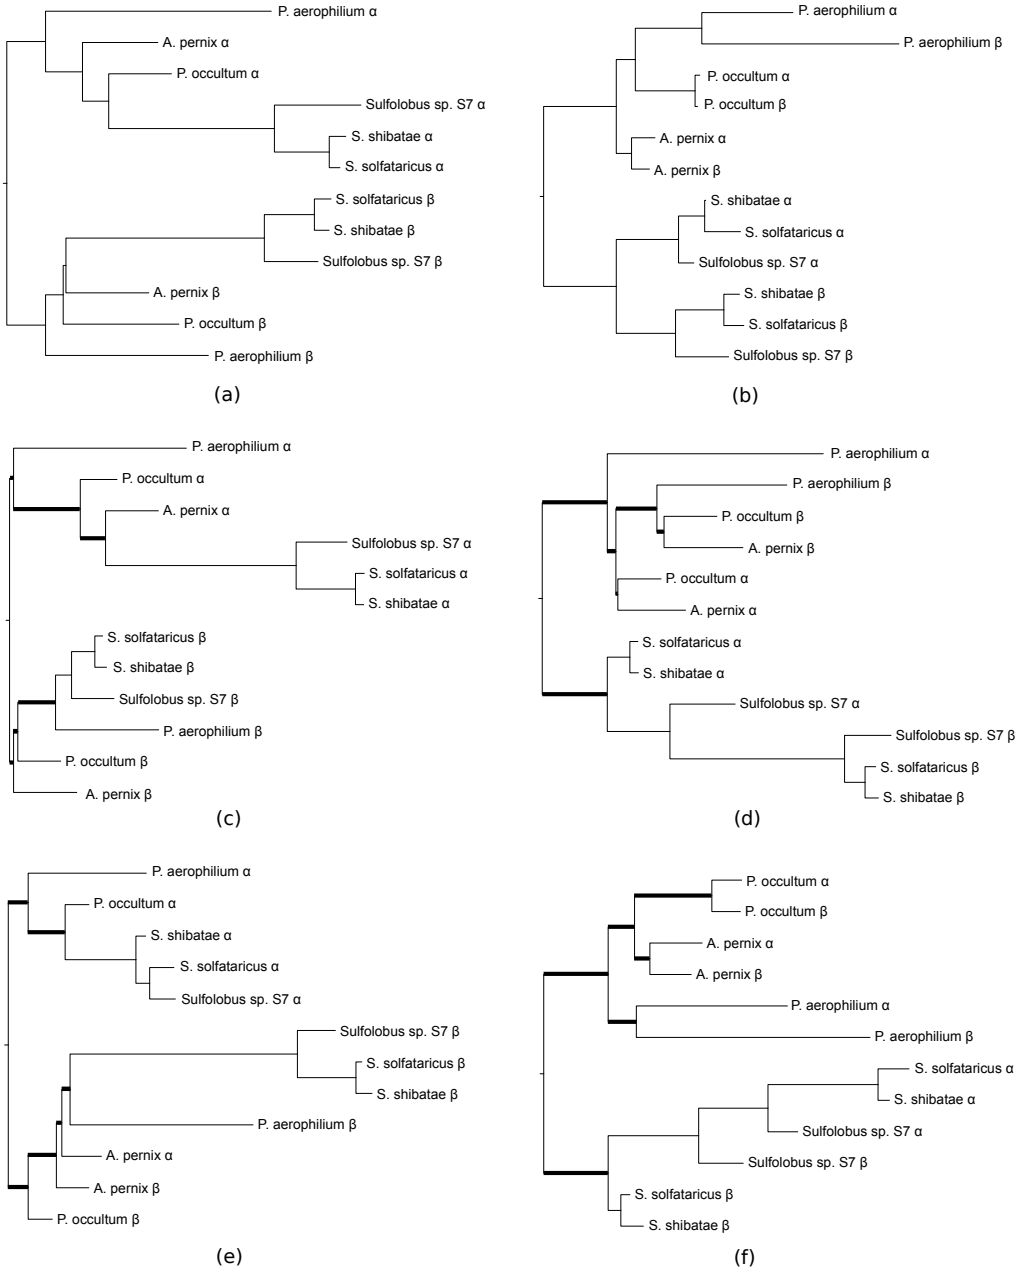

FIGURE 4. Phylogenies for the chaperonin data set. (a) The duplication topology, and (b) the conversion topology. The root splits of these two trees are referred to as the duplication and *Sulfolobus* splits in the main text. Leaves are labelled α or β according to the copy of the paralog. (c) and (d) show the phylogenies at the two ends of the principal geodesic for the sample of trees with window length 200. Edges have been thickened to illustrate the transition between the trees between the phylogenies along the principal geodesic: thickened edges in (c) are simultaneously shrunk to zero length to be replaced by the thickened edges in (d). This transition is followed by a change in topology within the *Sulfolobus* clade. (e) and (f) show the phylogenies at the two ends of the principal geodesic for the sample of trees from the apical domain with window length 100. Again, the thickened edges are simultaneously shrunk to zero length in (e) to be replaced by those in (f). There are separate changes in the topology of the *Sulfolobus* clade.

topologies in the original study, although there is some similarity. In particular, the two sequences corresponding to each of the species *P. occultum*, *A. pernix* and *P. aerophilum* are not grouped together in separate clades, as they are in the conversion topology. This is not entirely surprising, as the conversion topology was associated with a small proportion of sites in the original alignment, and variability in the phylogenies from the other sites will have played a proportionately higher role in construction of the principal geodesic. In addition, the conversion topology was associated with three separate contiguous regions each containing approximately 100 nucleotides, and so with a sliding window of length 200 the signal from these regions may have been obscured. A principal geodesic was therefore constructed for a sample of trees generated using a sliding window of length 100, with the windows restricted to cover a particular region, called the apical domain, identified by Archibald and Roger [23] as containing the loci affected by gene conversion. We call this the *apical* sample. This sample contained 50 trees, and the principal geodesic had $d_\perp^2 = 8.6$, $r_\gamma^2 = 26\%$. Ten runs of the algorithm converged to the same geodesic. Fig. 4 shows the trees corresponding to the ends of the principal geodesic. These have topologies very similar to the duplication and conversion topologies. The principal geodesic therefore identifies the main source of variability within the sample as coming from the support for these different topologies within the apical domain.

We computed the projection of the full sample of trees from the sliding window with length 200 onto the principal geodesic computed for the apical sample. The results are shown in Fig. 5. The relative position of the projection of a tree along the principal geodesic reflects its similarity to the two alternatives in Fig. 4 (e) and (f). The peak in the middle of the plot corresponds to a set of trees with the conversion topology. The graph acts as a 1-dimensional summary of the sample: the y-coordinate of each tree reflects particular features associated with the tree (for example trees above the dashed line are likely to contain the *Sulfolobus* split, and trees below, the duplication split). The original analysis of the chaperonin data set[23] includes a similar plot, based on a likelihood ratio, for testing relative support for the two conflicting topologies. The likelihood ratio approach for testing support is probably advantageous for assessing which regions of an alignment are associated with the two different topologies. However, it relies on prior knowledge of those topologies. In contrast, score plots like Fig. 5 can be prepared for principal geodesics directly from a sample of trees without such knowledge.

The score plot in Fig. 5 contains an apparent cluster of projected points at distance 0.16 along the geodesic which shows as a horizontal line on the plot. This point on the geodesic corresponds to a tree with

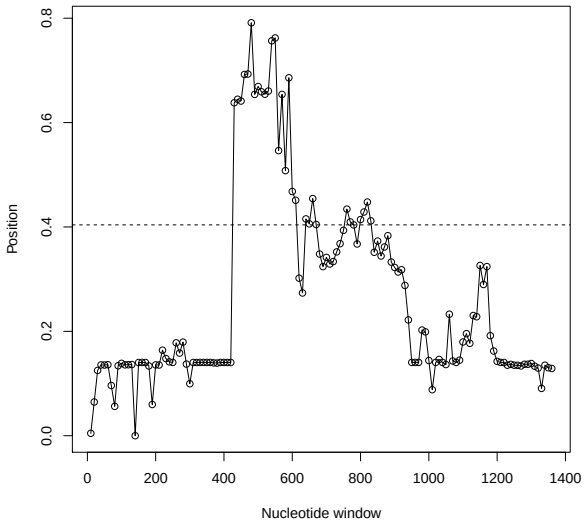

FIGURE 5. Projection of the full sample of trees for the chaperonin analysis onto the principal geodesic constructed from the apical sample . The x-coordinate is the starting point of each window in the alignment, and the y-axis is the position of the projected tree along the principal geodesic. Small y values correspond to trees near (e) in Fig. 4 while large values correspond to trees near (f). The apical region is approximately between $x = 550$ and $x = 1050$. The dashed line corresponds to the tree on the geodesic in which the thickened lines in Fig. 4(e) have been shrunk to zero.

the same topology as Fig. 4(e) but for which the three *Sulfolobus* α sequences are unresolved, and so it is associated with a change in the topology with the the *Sulfolobus* α clade. In order to investigate this effect, we prepared plots of the length of the three possible splits which resolve the *Sulfolobus* α clade versus window position, assigning length zero if the split in question was not present in a tree. These plots did not explain the pattern of scores in Fig. 5: specifically, no particular change in split lengths was observed at nucleotide positons 400 and 950. The full collection of splits and associated lengths for each tree must be taken into account to determine its projection, and it seems as though other features of the principal geodesic beyond these three particular splits determine where trees project relative to the change in the *Sulfolobus* α clade topology on the geodesic. In contrast to Euclidean space, the volume of the region which projects to a particular point on a geodesic segment varies from point to point in treespace, and this might be the cause of the apparent clustering of projected points.

	Full sample		Apical sample	
GeoPhytter	$d^2_\perp = 27.5$,	$r^2_\gamma = 43\%$	$d^2_\perp = 8.6$,	$r^2_\gamma = 26\%$
ΦPCA	$d^2_\perp = 39.7$,	$r^2_\gamma = 26\%$	$d^2_\perp = 12.8$,	$r^2_\gamma = 12\%$
Feragen *et al*	$d^2_\perp = 33.5$,	$r^2_\gamma = 32\%$	$d^2_\perp = 10.4$,	$r^2_\gamma = 15\%$

TABLE 1. Comparison of summary statistics for geodesics constructed for the full sample of trees and apical sample using the method presented in this paper (GeoPhytter) and extisting methods.

This has implications for principal geodesic analysis and dimensional reduction in treespace more generally, since it is desirable to obtain reductions in which distinct sampled trees are well discriminated.

Principal geodesics for the full sample and apical sample were also constructed using ΦPCA and the algorithm in [20]. The algorithm presented in [20] was run for $20,000$ iterations. This value was chosen since it gave a similar run-time to GeoPhytter. The results are summarized in Table 1. ΦPCA produced geodesics with a poor fit to the data. The geodesics constructed by GeoPhytter fit the data much better and do not lie in the class of 'simple' geodesics required by ΦPCA. For example, all the trees on the geodesic constructed by ΦPCA for the apical sample contained the *Sulfolobus* split, so the geodesic failed to represent the variability caused by gene conversion. As the table shows, the method of Feragen *et al* performs better than ΦPCA but GeoPhytter is substantially better than both methods in terms of smaller $d^2_\perp$ values and larger r^2_γ proportions.

3.2. **Parametric bootstrap data set.** In order to assess the performance of the algorithm on larger trees, we re-analysed a parametric bootstrap sample of trees considered in [19]. This sample was simulated from an underlying tree containing 41 taxa representing major eukaryote groups with an outgroup of archaea. The tree contained two long branches, corresponding to *microsporidia* and *guillardia*, in addition to the long branch leading to the archaeal outgroup. This simulated sample was analysed to explore the possibility of using principal geodesics to capture long branch attraction (LBA) effects. The data were transformed prior to construction of principal geodesics by scaling edges: edge lengths were scaled so that every split had unit mean length, with the mean for each split taken across trees containing that split in the data set. This transformation was performed so that variability was assessed as being relative to edge length, by scaling up the variability in short edges. Analysing transformed data in this way can be thought of as being analogous to using the sample correlation matrix instead of the covariance matrix in standard PCA. More details are given in

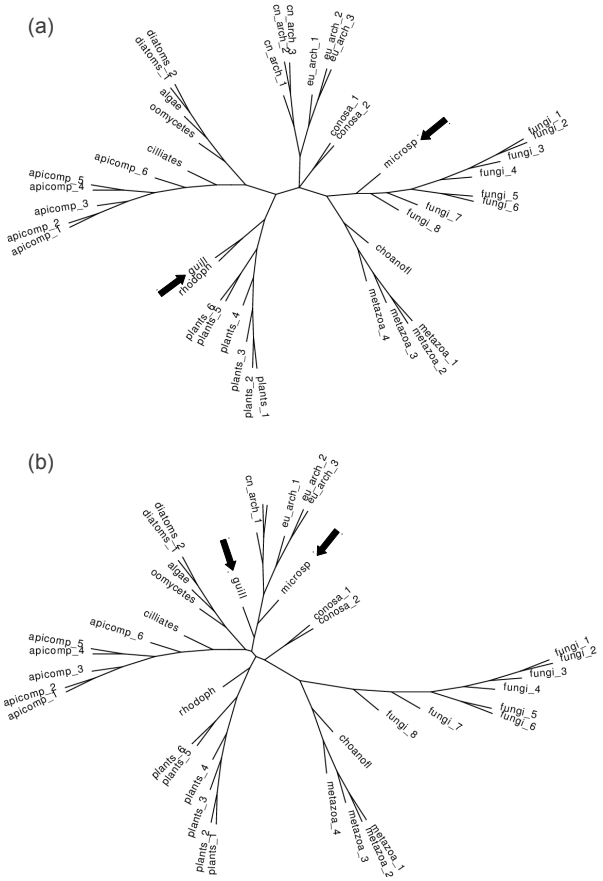

FIGURE 6. Trees forming the ends of the principal geodesic for the simulated bootstrap sample of trees. The arrows mark the taxa *guillardia* and *microsporidia* which have long pendant edges in the original phylogenies. The trees have been normalized so that each edge in the data set has unit mean length. *Guillardia* and *microsporidia* move around the tree as the principal geodesic is traversed, from their original positions in (a) to be grouped alongside the archaea (labelled 'cn arch' and 'eu arch') in (b).

[19]. The analysis was performed ignoring pendant edge lengths in the sample.

Fig. 6 shows the end-points of the principal geodesic constructed using the algorithm in Sec. 2.2. As the principal geodesic is traversed, the taxa with long branches, *microsporidia* and *guillardia*, both 'float' around the tree from their initial positions near fungi and plants respectively, to be grouped next to the archaea. The principal geodesic had $d_\perp^2 = 1290$ and $r_\gamma^2 = 13\%$ in comparison to the geodesic constructed

with ΦPCA which had $d_\perp^2 = 1507$ and $r_\gamma^2 = 10\%$. Both *microsporidia* and *guillardia* move around the tree, in contrast to the results from ΦPCA for which only one of the taxa moved in this way. The principal geodesic does not lie in the class of simple geodesics considered by ΦPCA, and so ΦPCA is not able to capture the same result. When viewed as an animation of trees, the principal geodesic gives an immediate visual representation of the LBA effect present in the data set.

The value $r_\gamma^2 = 13\%$ appears smaller than the values typically obtained in a standard principal components analysis. This is in part due to the high-dimensional nature of treespace, and the following analogy can be used to interpret r_γ^2. Principal geodesic analysis in treespace with N taxa can be compared with standard PCA in $\mathbb{R}^{N-3}$ when pendant edges are ignored. Consider the analysis of multivariate normal data for which the variance is σ^2 along the principal axis, and τ^2 in all other orthogonal directions. It follows that the ratio $\sigma^2 : \tau^2$ is $(N-4)r^2/(1-r^2)$. For $r^2 = 13\%$ and $N = 41$ as above this gives a value of 5.5. Under this Euclidean analogy, and assuming the other effects in the data consist of isotropic noise, the principal geodesic for the bootstrap sample has an associated variance which is a factor 5.5 times the noise variance. Of course, the analogy with Euclidean PCA is approximate, but nonetheless this calculation suggests that the LBA effect is significantly greater than random noise.

3.3. **Other data sets.** We constructed principal geodesics for several other data sets in order to assess the performance of the algorithm, and give brief details here to indicate the type of results that can be obtained via principal geodesic analysis.

Overall scale: Sometimes gene trees vary in the overall scale of the phylogenies. We re-analysed the metazoan data set considered in [19], and obtained a principal geodesic with $r_\gamma^2 = 57\%$. Although this principal geodesic represented some changes in topology, the main feature was the difference in total length of the trees at either end of the geodesic.

Single NNI: We analysed a well-known data set of 106 gene trees for 8 species of yeast [24], obtained by maximum likelihood inference. The principal geodesic passed through the majority consensus topology of the 106 gene trees and involved a single nearest-neighbor interchange between the subtrees ((*C.albicans*, *S.Kluyveri*), *S.castellii*) and (*C.albicans*, (*S.Kluyveri*, *S.castellii*)). The principal geodesic had $r_\gamma^2 = 63\%$, suggesting that the majority of variability in the data set was captured by this single NNI.

Isotropy: Some data sets are not well represented by a geodesic segment in tree space. In a recent study of turtle evolution [25], phylogenies were inferred for various species of turtle, but the individual representing each species was selected at random from a fixed pool of individuals. Phylogenies were constructed for 100 different sets of individuals sampled in this way. The principal geodesic constructed for this data set had r_γ^2 close to the isotropic baseline discussed in Sec. 2.6. Multiple runs of the algorithm failed to converge to a single geodesic: multiple local optima existed with similar values of the objective function. This suggests, by analogy with the Euclidean case, that the data are 'isotropic' in treespace with no distribution around a particular geodesic direction.

4. Discussion

4.1. **Computational issues.** Construction of principal geodesics in treespace is computationally demanding, in most part due to the huge number of possible topologies for the end-points of geodesic segments. Since our algorithm is stochastic it is not possible to give an overall algorithmic complexity, but each evaluation of the objective function (1) has order $O(N^4 \times n)$ where N is the number of taxa and n the number of data points. Calculation of the objective can be distributed across parallel processors simply by splitting up the data set into subsets, and so the algorithm speed increases linearly with the number of processors. The run times on a standard desktop computer were approximately 25 minutes for the chaperonin data set and 2 hours for the bootstrap sample of trees to obtain convergence from a single starting point. These times were obtained using 4 cores on a Intel Core i7 CPU running at 2.93GHz with 8Gb memory. However, very strict convergence criteria were used for the examples when preparing this article. In general the value of the objective function decreases approximately exponentially as the algorithm proceeds, and the majority of the runtime was spent performing minor adjustments to the segment γ close to the final value. Faster times could be obtained if the user accepted a greater degree of approximation. In practice, the algorithm is probably limited to work with samples containing fewer than 100 taxa and 500–1000 trees on a standard desktop computer. However, parallel computing resources would enable larger data sets, particularly in terms of the number of trees, to be analysed. The algorithm is implemented in java, and software is freely available from the web site given in the abstract. An important part of the software is a tool for visualizing geodesics in treespace. In a similar way to phylogenetic network diagrams, it enables users to relate different topologies associated with a sample. The user drags a slider to traverse the geodesic and view the corresponding smoothly changing phylogenetic tree. The collection of topologies

along a given geodesic segment is not always compatible with a single rotational ordering of taxa in the plane. Finding a rotational ordering for each topology in such a way as to minimize changes in ordering as the geodesic is traversed is problematic; the geodesic viewer uses simple heuristics to obtain a reasonable solution.

4.2. **Further research.** Summarizing samples and distributions on treespace is a fundamentally difficult problem due to the high dimensionality and non-Euclidean nature of treespace. Dimensional reduction, by fitting appropriate low-dimensional objects to samples of trees, is an obvious approach to adopt. This article has focussed on the construction of an analog of the first principal component of a sample, without consideration of the higher order components. Construction of the second order component, for example, requires an analogue of a plane or some other 2-dimensional surface in treespace, but the theory of higher-dimensional surfaces in treespace has not yet been developed. One possibility, a natural extension of the geodesic segments considered in this paper, is to seek a configuration of three points in treespace and consider the projection of the data onto the convex hull of these points. However, it is possible that completely different objects – not necessarily based on the geodesic geometry – might form better descriptors of distributions of phylogenies. Principal components analysis has been described in this article in terms of a least squares procedure rather than by reference to a probabilistic model. In this way, the sum of squared distances $d_\perp^2 = \sum d(x_i, P_\gamma(x_i))^2$ has essentially played the role of the likelihood of the data. The likelihood for multivariate normal models on $\mathbb{R}^k$ has exactly this form, and this leads to the description of standard PCA as a least squares procedure, as given in the introduction. We adopted the least squares approach in treespace by analogy, but a probabilistic model and a more fully developed theory of distributions on treespace would be highly desirable.

References

[1] D. Bryant, "A classification of consensus methods for phylogenetics," in *Bioconsensus*. American Mathematical Society, 2003, pp. 163–183.

[2] E. N. Adams, "Consensus techniques and the comparison of taxonomic trees," *Syst. Zool.*, vol. 21, no. 4, pp. 390–397, 1972.

[3] ——, "N-trees as nestings: complexity, similarity, and consensus," *J. Classif.*, vol. 3, no. 2, pp. 299–317, 1986.

[4] F. J. Lapointe and G. Cucumel, "The average consensus procedure: combination of weighted trees containing identical or overlapping sets of taxa," *Syst. Biol.*, vol. 46, no. 2, pp. 306–312, 1997.

[5] B. Holland and V. Moulton, "Consensus networks: a method for visualising incompatibilities in collections of trees," in *Algorithms in Bioinformatics*. Springer, 2003, lecture notes in computer science.

[6] C. Stockham, L. Wang, and T. Warnow, "Statistically based postprocessing of phylogenetic analysis by clustering," *Bioinformatics*, vol. 18, pp. S285–S293, 2002.
[7] J. Koperwas and K. Walczak, "Clustering of leaf-labelled trees," in *Adaptive and Natural Computing Algorithms*. Springer, 2007, pp. 702–710.
[8] D. Hillis, T. Heath, and K. St. John, "Analysis and visualization of tree space," *Syst. Biol.*, vol. 54, no. 3, pp. 471–482, 2005.
[9] L. Billera, S. Holmes, and K. Vogtmann, "Geometry of the space of phylogenetic trees," *Adv. Appl. Math.*, vol. 27, pp. 733–767, 2001.
[10] P. T. Fletcher, C. Lu, S. M. Pizer, and S. Joshi, "Principal geodesic analysis for the study of nonlinear statistics of shape," *IEEE T. Med. Imaging*, vol. 23, no. 8, pp. 995–1005, 2004.
[11] S. Huckemann and H. Ziezold, "Principal component analysis for riemannian manifolds, with an application to triangular shape spaces," *Adv. Appl. Probab.*, vol. 38, no. 2, pp. 299–319, 2006.
[12] P. T. Fletcher, C. Lu, and S. Joshi, "Statistics of shape via principal geodesic analysis on Lie groups," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2003, pp. 1–95.
[13] S. Huckemann and T. Hotz, "Principal component geodesics for planar shape spaces," *J. Multivariate Anal.*, vol. 100, no. 4, pp. 699–714, 2009.
[14] H. Wang and J.S. Marron, "Object oriented data analysis: sets of trees," *Ann. Stat.*, vol. 35, no. 5, pp. 1849–1873, 2007.
[15] J.S. Marron, S. Jung, and I. Dryden, "Speculation on the generality of the backward stepwise view of pca," in *Proceedings of the international conference on Multimedia information retrieval*. ACM, 2010, pp. 227–230.
[16] M. Bacak, "Computing medians and means in Hadamard spaces," *arXiv:1210.2145 [math.MG]*, 2012. [Online]. Available: http://arxiv.org/abs/arXiv:1210.2145
[17] E. Miller, M. Owen, and J. Provan, "Averaging metric phylogenetic trees," *arXiv:1211.7046 [math.MG]*, 2012. [Online]. Available: http://arxiv.org/abs/1211.7046
[18] K. Sturm, "Probability measures on metric spaces of nonpositive curvature," *Contemp. Math.*, vol. 338, pp. 357–390, 2003.
[19] T. Nye, "Principal components analysis in the space of phylogenetic trees," *Ann. Stat.*, vol. 39, pp. 2716–2739, 2011.
[20] A. Feragen, M. Owen, J. Petersen, M. Wille, L. Thomsen, A. Dirksen, and de Bruijne M., "Tree-space statistics and approximations for large-scale analysis of anatomical trees," in *23rd biennial International Conference on Information Processing in Medical Imaging (IPMI)*, 2013.
[21] M. Owen and J. Provan, "A fast algorithm for computing geodesic distances in tree space," *IEEE/ACM Trans. Comp. Biol. and Bioinf.*, vol. 8, no. 1, pp. 2–13, 2010.
[22] J. Spall, *Introduction to stochastic search and optimization: estimation, simulation, and control*. Wiley, 2003.
[23] J. Archibald and A. Roger, "Gene duplication and gene conversion shape the evolution of archaeal chaperonins," *J. Mol. Biol.*, vol. 316, no. 5, pp. 1041–1050, 2002.
[24] A. Rokas, B. Williams, N. King, and S. Caroll, "Genome-scale approaches to resolving incongruence in molecular phylogenies," *Nature*, vol. 425, pp. 798–804, 2003.

[25] P. Spinks, R. Thomson, G. Pauly, C. E. Newman, G. Mount, and H. Shaffer, "Misleading phylogenetic inferences based on single-exemplar sampling in the turtle genus pseudemys." *Mol. Phylogenet. Evol.*, vol. 68, pp.

E-mail address: tom.nye@ncl.ac.uk

School of Mathematics and Statistics, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK